

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

PCT

NOTIFICATION OF ELECTION

Date of mailing: 30 March 2000 (30.03.00)	in its capacity as elected Office
International application No.: PCT/GB99/03114	Applicant's or agent's file reference: A25570 WO
International filing date: 17 September 1999 (17.09.99)	Priority date: 18 September 1998 (18.09.98)
Applicant: EVANS, Paul, Andrew et al	

1. The designated Office is hereby notified of its election made:

in the demand filed with the International preliminary Examining Authority on:

17 February 2000 (17.02.00)

in a notice effecting later election filed with the International Bureau on:

2. The election was

was not

made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

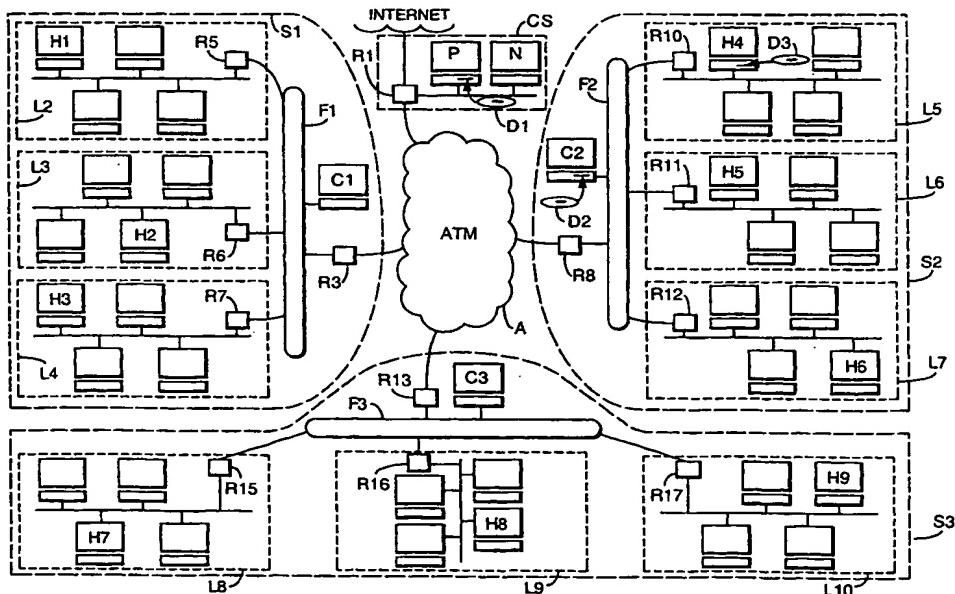
<p>The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland</p> <p>Facsimile No.: (41-22) 740.14.35</p>	<p>Authorized officer:</p> <p style="text-align: center;">J. Zahra</p> <p>Telephone No.: (41-22) 338.83.38</p>
---	--



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 12/18	A1	(11) International Publication Number: WO 00/18068 (43) International Publication Date: 30 March 2000 (30.03.00)
<p>(21) International Application Number: PCT/GB99/03114 (22) International Filing Date: 17 September 1999 (17.09.99)</p> <p>(30) Priority Data: 98307623.3 18 September 1998 (18.09.98) EP</p> <p>(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only): EVANS, Paul, Andrew [GB/GB]; Flat 3, 54 Anglesea Road, Ipswich, Suffolk IP1 3PW (GB). BRISCOE, Robert, John [GB/GB]; Home Farm, Parham, Woodbridge, Suffolk IP13 9NW (GB). RUDKIN, Steven [GB/GB]; 52 Corder Road, Ipswich, Suffolk IP4 2XD (GB).</p> <p>(74) Agent: NASH, Roger, William; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).</p>		(81) Designated States: AU, CA, JP, SG, US, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>		

(54) Title: HIERARCHICAL MULTICASTING



(57) Abstract

A packet network and method of operating the same is described. The network is arranged into a hierarchy, having one organisation-level portion (CS, P, R2) and three site-level portions (S1, S2, S3). By using forwarding computers and assigning multicast addresses in respect of each portion, the total number of multicast addresses required to send multicast messages to all possible subsets of potential recipients is reduced. This alleviates routers in the network of the burden associated with storing and processing a large amount of multicast addresses.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		

HIERARCHICAL MULTICASTING

The present invention relates to a method of operating a transmitter to multicast data blocks over a network.

5

In network terminology, sending a data block (i.e. a message, packet, frame or the like) across a network to a single recipient is known as unicasting. Sending a data block to all receivers connected to a network is known as broadcasting. Sending a data block to a set of recipients that form a subset of all receivers connected to a 10 network is referred to as multicasting. Such a set of recipients is known as a multicast group.

In broadcast networks (e.g. nearly all Local Area Networks), all receivers receive any data block placed the network. Each receiver has a network interface card 15 that identifies frames that are intended for that receiver. In order to identify which multicast data blocks are intended for that receiver, the network interface card must store an indication of each multicast address for which the receiver is a recipient.

20 In networks comprising switching nodes interconnected by point-to-point links, each switching node must store, for each multicast group, an indication of which of the links connected to it lead toward members of the multicast group.

25 In both types of networks therefore, the amount of memory required in network components increases with the number of multicast addresses in use.

Conventional use of multicast networks often results in the use of more multicast addresses than is necessary. Furthermore, the amount of traffic generated in passing one or more multicast messages across the network is often unnecessarily 30 high.

Although the use of multicasting is currently at a level where these problems do not impact heavily on network performance, a rapid rise in the use of multicast is

likely to take place over the next few years. This means that there is a need to provide a more efficient way of using networks that provide a multicast capability.

According to a first aspect of the present invention, there is provided a method of
5 operating a transmitter of data blocks connected via a network having a multicast capability to a plurality of receivers of said data blocks, said transmitter having access to one or more directories listing sets of recipients selected from said receivers and corresponding multicast addresses; said method comprising:

- 10 a) finding the set of recipients to which a data block is to be sent;
 - b) examining said one or more directories to find one or more respective multicast addresses corresponding to said set of recipients;
 - 15 c) addressing one or more data blocks to said one or more respective multicast addresses; and
 - d) transmitting said one or more data blocks over said network.
- 20 By having one or more directories that list possible selections of recipients and corresponding multicast addresses, multicast data blocks can be addressed based on the set of recipients to which it is found the data block is to be sent. This leads to the process of multicasting a data block being more easily adaptable to the requirements of different users.

25

For groups with only one member, the directories can provide a unicast address instead of a multicast address.

- To give an example, in some embodiments of the present invention, said finding
30 step comprises:
 - a) receiving one or more indications that an earlier data block addressed to a selected set of recipients was not successfully received by one or more of said recipients; and

b) analysing said indications to find the subset of said selected set of recipients that did not successfully receive said earlier data block; wherein said examining step involves examining said one or more directories to find the multicast address corresponding to said subset of recipients.

5

By finding the set of recipients that require re-transmission of the data block, and then finding from the one or more directories the multicast address that corresponds to that set of recipients, the re-transmission can easily be directed only towards those recipients that failed to receive the earlier data block. This 10 results in less load being placed on the network than would occur in accordance with conventional methods where the data block is re-transmitted to the same multicast address as before once one or more recipients have indicated that they have not received the data block.

15 To give another example of the improved adaptability provided by operating the transmitter in accordance with the present invention, in some embodiments said finding step involves:

a) determining that a general data block is to be sent to a selected plurality of sets of recipients; and
20 b) unifying said plurality of sets of recipients to find a unified set of recipients; wherein said examining step involves examining said one or more directories to find from said address data the multicast address which corresponds to said unified set.

25 In this case only one multicast message need be sent across the network, as opposed to the plurality of multicast messages that would have to be sent using a conventional method.

It is to be expected that 'push' news services, Publish and Subscribe services and 30 the like will soon migrate to the use of multicasting. 'Push' news services normally allow a user to select subjects of interest. One way of achieving the bandwidth savings offered by the use of multicast (whilst allowing a user to be selective as to which subjects he received) would be to allocate a multicast address to each subject. However, given the large number of possible subjects

this method would result in the components of the network having to handle a large number of multicast addresses.

- To address this problem, in a further embodiment of the present invention, said
- 5 transmitter further has access to type data listing data block type identifiers and corresponding sets of recipients, and said finding step involves:
- a) finding a type identifier associated with said data block; and
 - b) examining said type data to find the set of recipients associated with said type identifier.

10

The transmitter might update the type data in response to a request from a recipient to be sent data blocks containing a given type identifier, or in response to a request not to receive any further data blocks containing that type identifier.

- 15 The data block type identifier might, for example, indicate to which subject the information carried in the data block relates. In conventional methods of operating a transmitter to use a multicast-capable network, a one-to-one subject-to-multicast address mapping is used. However, it will be seen that the above embodiment enables a many-to-one subject-to-multicast address mapping to be used instead.
- 20 In networks where the number of selections of recipients is less than the number of subjects, the number of multicast addresses that must be handled by the network is reduced.

It is even possible to utilise the present invention to reduce the number of

25 multicast addresses that must be handled by a network where the number of possible sets of recipients is very high (it will be realised that the number of sets of recipients grows exponentially with the number of receivers).

One way of achieving this is to divide the recipients into groups and provide a

30 separate directory for each group. The number of multicast addresses required then grows exponentially with the number of receivers in each group, rather than growing exponentially with the number of receivers connected to the network.

Another way of achieving this is to provide, in accordance with a second aspect of the present invention, a method of operating an internetwork, said internetwork comprising:

- a first level transmitter of data blocks connected via a first level network to a plurality of first level network receivers which include a plurality of second level transmitters in turn connected via respective second networks to respective pluralities of second level network receivers,
- 5 said second level transmitters having access to respective second level network directories, each second level network directory listing data block type identifiers and corresponding multicast addresses for the second level network to which the directory relates, said method comprising the steps of:
 - operating said first transmitter to transmit data blocks over said first network, and further operating said first transmitter to include a data block type identifier in each data block;
 - 10 operating each of said second level transmitters to forward data blocks over the corresponding second level network by:
 - a) extracting said type identifier from a data block received from the first transmitter;
 - b) examining the appropriate second level network directory to find the second level network multicast address corresponding to said type identifier;
 - 20 c) addressing said data block to said second level network multicast address; and
 - d) transmitting said data block over said second level network.

The term internetwork is used here to mean a plurality of interconnected networks (which themselves therefore form a larger network). Here the networks are arranged into a hierarchy, the connections between the first and second level networks being made by respective second level transmitters. It will be realised that the second level network directories are similar to the group directories mentioned above and result in a corresponding reduction in the number of multicast addresses that must be handled in the larger network.

In hierarchical embodiments of the present invention one or more of the smaller networks can be operated in accordance with one or more of the above embodiments.

In embodiments of the second aspect of the present invention, the transmitter at one level of the hierarchy (upper transmitter) places a data block type identifier in each data block that it transmits – the transmitter at the next level down in the 5 hierarchy (the lower transmitter) is provided with access to a directory which sets out the multicast addresses associated with each data block type identifier. On receiving a data block from the upper transmitter, the lower transmitter reads the data block type identifier, finds the multicast address associated with that data block type identifier, re-addresses the data block accordingly, and forwards the 10 data block across its network. The hierarchy can have any number of levels.

In a hierarchical arrangement a request from a recipient to be sent data blocks of a particular type need only travel up the hierarchy as far as the intermediate transmitter that already receives data blocks of that type. This advantageously 15 reduces network traffic.

Those skilled in the art will realise that where the smaller networks operate in accordance with the Internet Protocol suite, administrative scoping (a technique which limits the nodes through which a multicast message may pass) can be used 20 to enable the re-use of multicast addresses in separate smaller networks, thereby decreasing further the number of multicast addresses that need be handled by the larger network.

By having the transmitters intermediate the smaller networks cache data blocks for 25 a predetermined length of time, re-transmission of the data block can be limited to the smaller network in which the transmission failure occurred. Similarly, flow-control can be restricted to a smaller network that is experiencing congestion.

It will be realised that transmitters at all levels of the hierarchy can be provided 30 with a directory listing multicast addresses for the smaller networks they transmit data blocks over.

Thus, according to a third aspect of the present invention, there is provided a method of operating a packet network to transmit a plurality of packets to

respective different subsets of possible recipients, said method comprising the steps of, for one or more portions of the network, assigning a common local multicast address to packets which are destined for different subsets of possible recipients but which share a subset of forwarding nodes within said portion.

5

According to a fourth aspect of the present invention, there is provided a data block forwarding apparatus operable to forward data blocks received from a first network over a second network to a set of second network recipients selected from a plurality of second network receivers, said apparatus comprising:

- 10 an input connectable to said first network to receive data blocks containing a data block type identifier;
- an output connectable to said second network;
- one or more processors;
- a forwarding database store storing data representing data block type identifiers and corresponding sets of recipients;
- 15 a directory store storing data representing sets of recipients and corresponding multicast addresses;
- a program store storing instructions executable by said one or more processors to:
- 20 a) read said data block type identifier from a data block received from said first network;
- b) examine said forwarding database store to find the set of recipients corresponding to said data block type identifier;
- c) examine said directory store to find the multicast address associated with the
- 25 set of recipients found in step b);
- d) re-address said data block to the multicast address found in step c); and
- e) forward the data block via said output over said second network.

Such apparatus provides an transmitter intermediate two levels of the hierarchy
30 mentioned above.

There now follows, by way of example only, a description of specific embodiments of the present invention. The description is to be read in conjunction with the accompanying drawings, in which:

Figure 1 is a schematic view of an intranet;

Figures 2A and 2B show multicast address allocations for subsections of the
5 intranet;

Figure 3 shows the format of an application layer packet assembled in accordance
with one embodiment of the present invention;

10 Figures 4A, 4B and 4C illustrate the building up of subject-to-address mappings
used in the embodiment; and

Figures 5A, 5B, 5C and 6 illustrate more developed subject-to-address mappings.

15 A corporate intranet (Figure 1) comprises a central site area network (CS) which
includes an IEEE 802.3 local area network L1 to which a sports news computer N,
an organisation-level application packet forwarding computer P and a gateway
router R1 are connected. The gateway router R1 is also connected to the Internet
and an Asynchronous Transfer Mode (ATM) network A.

20

The ATM network A connects the gateway router R1 to three site area ingress
routers (R3, R8, R13). Each of the three site area ingress routers (R3, R8, R13) is
also connected to a respective one of three site-level internetworks. Each site
level internetwork comprises a Fibre Distributed Data Interface (FDDI) backbone

25 local area network (F1,F2,F3) which is connected by three building routers (R5 to
R7; R10 to R12; R15 to R17) to three respective IEEE 802.3 building local area
networks (L2 to L4, L5 to L7 and L8 to L10). Also connected to the FDDI network
(F1 to F3) is an application layer packet forwarding computer (C1, C2, C3).

30 Each building area network (L2 to L4, L5 to L7 and L8 to L10) is also connected to
a number of personal computers (PCs).

The sports news computer N at the organisation's central site includes a database containing a number of news articles about various sports. These news articles are updated every few hours.

- 5 One PC (H1 to H9) in each of the building area networks (L2 to L10) is designated as a news reader PC.

In configuring the intranet, each of the application layer packet forwarding computers (C1, C2, C3) is manually configured with a site area multicast directory

- 10 (Figure 2A). The directory lists possible sets of recipients (left-hand column) and corresponding IEEE 802 addresses (right-hand column). IEEE 802 addresses can be forwarded over both the FDDI network (F1 to F3) and the building area IEEE 802.3 local area networks (L2 to L10). Those skilled in the art will realise that the building routers (R5 to R7; R10 to R12; R15 to R17) are therefore acting as
15 bridges in the present embodiment.

Further configuration involves the network interface cards of the news reader hosts H1 to H9 being configured to ensure that frames having addresses which relate to sets of recipients that include that host are passed from the network

- 20 interface card to the host itself. For example (see Figure 2A), the network interface card of news reader host H1 which is present in the sets (H1,H2,H3), (H1,H2), (H1,H3) and (H1) is configured to forward frames addressed to IEEE 802 addresses C000:0000:0001, C000:0000:0002, C000:0000:0003 (which those skilled in the art will recognise as multicast addresses) and also unicast address
25 0400:0000:0001.

Similar procedures are carried out in relation to the other site area internetworks (S2, S3).

- 30 The next stage in the configuration procedure is to supply a set-to-IP address conversion table to the organisation level packet forwarding computer P. This table is similar to that (Figure 2A) supplied to each of the site-level forwarding computers (C1, C2, C3), but the sets in this table are sets of site-level forwarding computers (C1,C2,C3) rather than sets of news reader PCs (H1 to H9). Also IP

addresses are used since the ATM network A is unable to forward IEEE 802.3 frames. One possible example of the contents of the set-to-IP address conversion table stored in the organisation-level forwarding computer P is shown in Figure 2B.

- 5 Further, each site area ingress router (R3,R8,R13) is configured to unicast packets having an appropriate one of the addresses in Figure 2B to the site-level forwarding computer (C1,C2,C3). For example, site area ingress router R8 is configured to send packets having one of the first, second, fourth or fifth addresses shown in Figure 2B to the site-level forwarding computer C2.

10

Those skilled in the art will realise that this will result in gateway router R1 having four news application-related multicast addresses in its multicast routing table. Each entry will include a mapping from the IP address to a corresponding VCI/VPI (Virtual Path Identifier/ Virtual Circuit Identifier) pair.

15

Once this configuration has taken place, users of the news reader PCs (H1 to H9) can subscribe to news articles about their chosen sport. Each news reader PC is controlled by a news reader program (installed on the computer from compact disc D3) which allows a user to select (using a Graphical User Interface for example) 20 sports for which they would like to receive news articles held in the news computer N. On such a request being made the program controls the PC to send a request packet (Figure 3) for news articles concerning the chosen sport to site-level forwarding computer (C1, C2, C3).

25 In the composition of the request packet, the news reader program controls the computer to provide data indicating that a file is sought 5, an indication that the files relates to a news application 10 and an identification of the subject 20 about which the news article is sought. The computer then operates in a conventional manner to successively add to the data: a UDP header 30, an IP header 40 and an 30 IEEE 802 header 60 and trailer 70.

On receiving the request, the site-level application packet forwarding computer (C1, C2, C3) is controlled by a forwarding table updating program (installed on that computer from compact disc D2) to establish which of the on-site news reader PCs

(H1 to H9) are now interested in news articles about the chosen sport. If the request is not already being met then the forwarding table is updated accordingly. Those skilled in the art will be able to generate suitable table updating program.

- 5 Were news reader PC H1 to send, for example, a 'subscribe' request for news articles concerning the subject 'pool', then the site-level forwarding computer operating under control of the program would update the forwarding table as shown in Figure 4A. It will be seen that the subject 'pool' is associated with H1's unicast IEEE 802 address.

10

Once the table has been updated, the packet forwarding program controls the site-level application packet forwarding computer (C1, C2, C3) to check to see whether the forwarding computer (C1, C2, C3) already subscribes to the requested subject. If it does not then the site-level forwarding computer is further controlled to send a

- 15 request to the organisation-level forwarding computer P.

On receipt of that request, the organisation-level forwarding computer P is controlled by a program similar to that described above in relation to the site-level forwarding computer (installed on the computer P from compact disc D1) to

20 establish which set of the site-level forwarding computers (C1,C2,C3) now requires news articles concerning the selected subject to be sent. The organisation-level forwarding computer is then further controlled to update its forwarding table accordingly. Again programs to cause this operation can be easily generated by those skilled in the art.

25

To continue with the earlier example, in response to the request from site-level application packet forwarding computer C1 for 'pool' articles, the organisation-level computer will update its forwarding table as shown in Figure 4B. It will be seen that the forwarding table indicates that news articles concerning 'pool' are to

30 be unicast to the IP address of site-level forwarding computer C1.

Using similar procedures to those described above, a request for articles relating to 'pool' might later be received by site-level forwarding computer C1 from news reader PC H2. The site-level packet forwarding computer C1 is therefore

controlled to amend its forwarding table to reflect the fact that both news reader H1 and news reader H2 should be sent articles concerning 'pool'. The site-level forwarding computer will therefore update its forwarding table as shown in Figure 4C.

5

- It will be seen that as two news reader PCs (H1, H2) have now requested articles relating to a subject, the IEEE 802 address associated with the subject is a multicast address (in the notation generally used, IEEE 802 multicast addresses begin with the hexadecimal number C000 and IEEE 802 unicast addresses begin 10 with hexadecimal number 0400).

It will be realised by those skilled in the art that 'unsubscribe' requests could also easily be implemented.

- 15 After a number of requests have been received from the various news reader PCs (H1 to H9) in the intranet, the forwarding tables will associate several subjects with each multicast or unicast address.

For example, the forwarding tables of site-level forwarding computers C1, C2 and 20 C3 might be as shown in Figures 5A, 5B and 5C respectively. At the same time, the forwarding table associated with organisation-level forwarding computer P might be as shown in Figure 6.

On receiving a news article, each of the site-level forwarding computers 25 (C1,C2,C3) is controlled by the forwarding program to extract both the news article (Figure 3, 50) and subject 20 from the payload of the incoming UDP protocol data unit and send another UDP protocol data unit containing the news article to the IEEE 802 address associated with the subject 20 in the table. Programs executable to control the forwarding computers to behave in this manner 30 can be generated by those skilled in the art.

To give an example, were the organisation-level computer P to receive a new news article concerning 'golf' and were the forwarding tables to be as shown in Figures 5A, 5B, 5C and 6 then the organisation-level forwarding computer P would be

- controlled to send the packet to IP address 229:274:1:27 (the multicast address corresponding to all three site-level forwarding computers (C1,C2,C3)). On receiving that packet the router R1 would fragment the IP packet and, in accordance with the routing table entries made at the configuration stage, send
- 5 the fragments as a stream of ATM cells having a corresponding VPI/VCI pair as mentioned above. Once the packet arrives at the second site-level forwarding computer C2, that forwarding computer C2 is controlled to reassemble the IP packet, extract the subject information (golf) from its payload and forward the IP packet in an appropriately addressed IEEE 802 frame. The third line of Figure 5B
- 10 shows the associated IEEE 802 address. Owing to the earlier configuration stage, the network interface cards of the hosts H4 and H6 pass the packet to the appropriate news reader hosts (H4 and H6 in this case) and then process the packet to present the news item to the user.
- 15 By providing the forwarding computers with further programs, they could be rendered operable to allow efficient 'cross-posting'. 'Cross-posting' involves the sending of a message which concerns a plurality of subjects. The subject field of a forwarded packet might contain an identification of a pair of subjects, for example. It can be envisaged that the site-level forwarding computers (C1,C2,C3) could
- 20 operate under control of the further program to identify the union of the sets of recipients for the two subjects and then forward the news item to the interested hosts. For example, assume the sets of recipients interested in various subjects are as shown in Figures 5A, 5B, 5C and 6, and that forwarding computer C1 has received a news item that relates to both hockey and golf. The forwarding
- 25 computer can easily be programmed to determine that both H1 and H3 are interested in the news item and thereafter send the news item to the IEEE 802 address that corresponds to the group (H1,H3) namely C000:0000:0003.

Similarly, in relation to reliable multicast protocol, a further program could be used

30 to control the forwarding computers to re-send messages only to the subset of recipients who did not acknowledge the message. Again, using the techniques of the above embodiment, the selection of an appropriate multicast/unicast address would be trivial.

In the above embodiment, the configuration of the network interface cards to pass appropriately addressed multicast frames to the news reader hosts H₁ to H₉ themselves is carried out manually. This configuration could be done automatically by having the site-level forwarding computers (C₁, C₂, C₃) unicast configuration messages to the news reader hosts (unicast addresses are available from the configuration tables (e.g. the last three lines in Figure 2A).

An alternative embodiment of the present invention does not have the site-level forwarding computers (C₁, C₂, C₃). At the configuration stage the organisation-level forwarding computer P is provided with a directory that lists multicast addresses for each of the 500 or so possible sets of news reader PC recipients. It will be realised that such an embodiment requires the central router R₁ to have routing table entries for all those addresses. However this still represents an improvement over conventional methods of using multicast which involve a multicast address being provided in relation to each subject to which a news item relates. The USENET service for example provides information on about 27,000 subjects. Hence, were a similarly comprehensive news service to be provided over the corporate intranet (Figure 1) using conventional techniques the central router would require 27,000 entries in its multicast routing table. It will be seen how the use of a multicast directory enables this to be reduced to 500 multicast routing entries in this alternative embodiment.

However, the first-described embodiment reduces the number of multicast routing entries in the gateway router R₁ even further. By arranging the corporate intranet as a hierarchy, with the ATM network A providing the upper level network and the site internetworks (S₁, S₂, S₃) providing the lower level networks, the central router need only have as many multicast forwarding table entries as there are selections of lower level networks (in the above described embodiment only four multicast entries and three unicast entries need be stored in central router R₁).

30

Hence, it will be seen how the use of a directory can reduce the burden on the intranet and how the use of multiple directories can reduce the burden on the network still further.

The above advantages are amplified in embodiments that have more than two hierarchical layers in the forwarding structure. For example, an embodiment can be envisaged in which town-level and country-level forwarding computers are used.

5

The above embodiment only enables the transmission of news articles from the news computer N to the news reader PCs (H1 to H9). A 'publish and subscribe' service could be provided by programming the news reader PCs (H1 to H9) to be operable to unicast messages containing news articles input by their users to the

10 news computer N. These messages could then be sent back down the hierarchy as described above.

A yet further embodiment of the present invention is similar to the alternative embodiment mentioned above in that the site-level forwarding computers 15 (C1,C2,C3) are not required. Nevertheless the number of multicast addresses handled by the network can be reduced without arranging the network into a hierarchy. To do this the central forwarding computer P is provided with a separate directory for each of the internetworks (S1,S2,S3). It will be realised that each of those directories will only require four multicast addresses.

20

More complex arrangements which are effective to obviate the need to send messages to the news computer at the top of the hierarchy might also be used.

In other embodiments many different applications might be supported. For 25 example, by replacing the subject field of the above embodiment with a field that indicated both an application and a subject, the embodiment could support a plurality of subject-based applications. One such application is the announcement of multicast sessions. Session announcement could include a subject identifier similar to that currently seen in Network News applications (e.g. 30 comp.internet.ietf.conference). It will be clear how the above embodiment might be altered to allow a user to selectively receive only those announcements which concern subjects in which he is interested. Such a facility could be provided instead of or in addition to the 'push' news service described above.

CLAIMS

1. A method of operating a transmitter of data blocks connected via a
5 multicast-capable network to a plurality of receivers of said data blocks, said
transmitter having access to one or more directories listing sets of recipients
selected from said receivers and corresponding multicast addresses; said method
comprising:
 - 10 a) finding the set of recipients to which a data block is to be sent;
 - b) examining said one or more directories to find respective one or more multicast
addresses corresponding to the set of recipients found in step a);
 - 15 c) addressing one or more data blocks to the respective multicast addresses found
in step b); and
 - d) transmitting said one or more data blocks over said network.
- 20 2. A method according to claim 1 wherein:
said finding step comprises:
 - a) receiving one or more indications that an earlier data block addressed to a
selected set of recipients was not successfully received by one or more of said
25 recipients; and
 - b) analysing said indications to find the subset of said selected set of recipients
that did not successfully receive said earlier data block;
said examining step involves examining said one or more directories to find the one
or more multicast addresses corresponding to said subset of recipients;
- 30 3. A method according to claim 1 wherein:
 - whereby said transmitting step involves re-sending said earlier data block to said
subset of recipients.

said finding step involves:

- a) determining that a general data block is to be sent to a selected plurality of sets of recipients; and
 - b) unifying said plurality of sets of recipients to find a unified set of recipients;
- 5 wherein said examining step involves examining said one or more directories to find the one or more multicast addresses which correspond to said unified set.

4. A method according to claim 1 wherein said transmitter further has access
10 to type data listing data block type identifiers and corresponding sets of recipients,
wherein said finding step involves:

- a) finding a type identifier associated with said data block; and
- b) examining said type data to find the set of recipients associated with said type identifier.

15

5. A method according to claim 4 wherein said type identifier is a subject-matter identifier indicating the subject-matter to which the data in the data block relates.

20 6. A method according to claim 4 or 5 wherein said type identifier finding step involves extracting a type identifier from a data block received at said transmitter.

7. A method according to claim 1 wherein said transmitter has access to a
25 plurality of group directories for respective groups of receivers.

8. A method of operating an internetwork, said internetwork comprising:
a first level transmitter of data blocks connected via a first level network to a plurality of first level network receivers which include a plurality of second level
30 transmitters in turn connected via respective second networks to respective pluralities of second level network receivers,
said second level transmitters having access to respective second level network directories, each second level network directory listing data block type identifiers

and corresponding multicast addresses for the second level network to which the directory relates, said method comprising the steps of:

- operating said first transmitter to transmit data blocks over said first network, and further operating said first transmitter to include a data block type identifier in each
5 data block;

operating each of said second level transmitters to forward data blocks over the corresponding second level network by:

- a) extracting said type identifier from a data block received from the first transmitter;
- 10 b) examining the appropriate second level network directory to find the second level network multicast address corresponding to said type identifier;
- c) addressing said data block to said second level network multicast address; and
- d) transmitting said data block over said second level network.

- 15 9. A method according to claim 8 wherein said second level transmitters cache said data blocks and operate to re-transmit data blocks over said second level network in accordance with the method of claim 2.

10. A method according to claim 8 wherein said first level transmitter has
20 access to a first level network directory listing sets of recipients and corresponding multicast addresses for the first level network and operates in accordance with the method of any one of claims 1 to 6.

11. A method according to any preceding claim wherein the format of said
25 multicast address is in accordance with the Internet Protocol suite.

12. A data block forwarding apparatus operable to forward data blocks received from a first network over a second network to a set of second network recipients selected from a plurality of second network receivers, said apparatus
30 comprising:

an input connectable to said first network to receive data blocks containing a data block type identifier;

an output connectable to said second network;

one or more processors;

a forwarding database store storing data representing data block type identifiers and corresponding sets of recipients;

a directory store storing data representing sets of recipients and corresponding multicast addresses;

5 a program store storing instructions executable by said one or more processors to:

a) read said data block type identifier from a data block received from said first network;

b) examine said forwarding database store to find the set of recipients 10 corresponding to said data block type identifier;

c) examine said directory store to find the multicast address associated with the set of recipients found in step b);

d) re-address said data block to the multicast address found in step c); and

e) forward the data block via said output over said second network.

15

13. A program storage device readable by a processing apparatus, said device embodying a program of instructions executable by the processing apparatus to perform method steps for transmitting a data block over a network to a set of recipients selected from a plurality of network receivers, said method steps 20 comprising steps according to any one of claims 1 to 6.

14. A computer program comprising computer program code adapted to perform the method steps of any one of claims 1 to 7 when said program is executed by a computer.

25

15. A method of operating a packet network to transmit a plurality of packets to respective different subsets of possible recipients, said method comprising the steps of, for one or more portions of the network, assigning a common local multicast address to packets which are destined for different subsets of possible 30 recipients but which share a subset of forwarding nodes within said portion.

16. A method according to claim 15 wherein said assignment step involves:

extracting, from the payload of each packet, destination data identifying the subset of possible recipients to which the packet is to be sent; and

deriving an associated local multicast address on the basis of said destination data; and

transmitting said packet to said associated local multicast address.

- 5 17. A method according to claim 16 wherein said destination data comprises subject identifying data.

17

丁
正

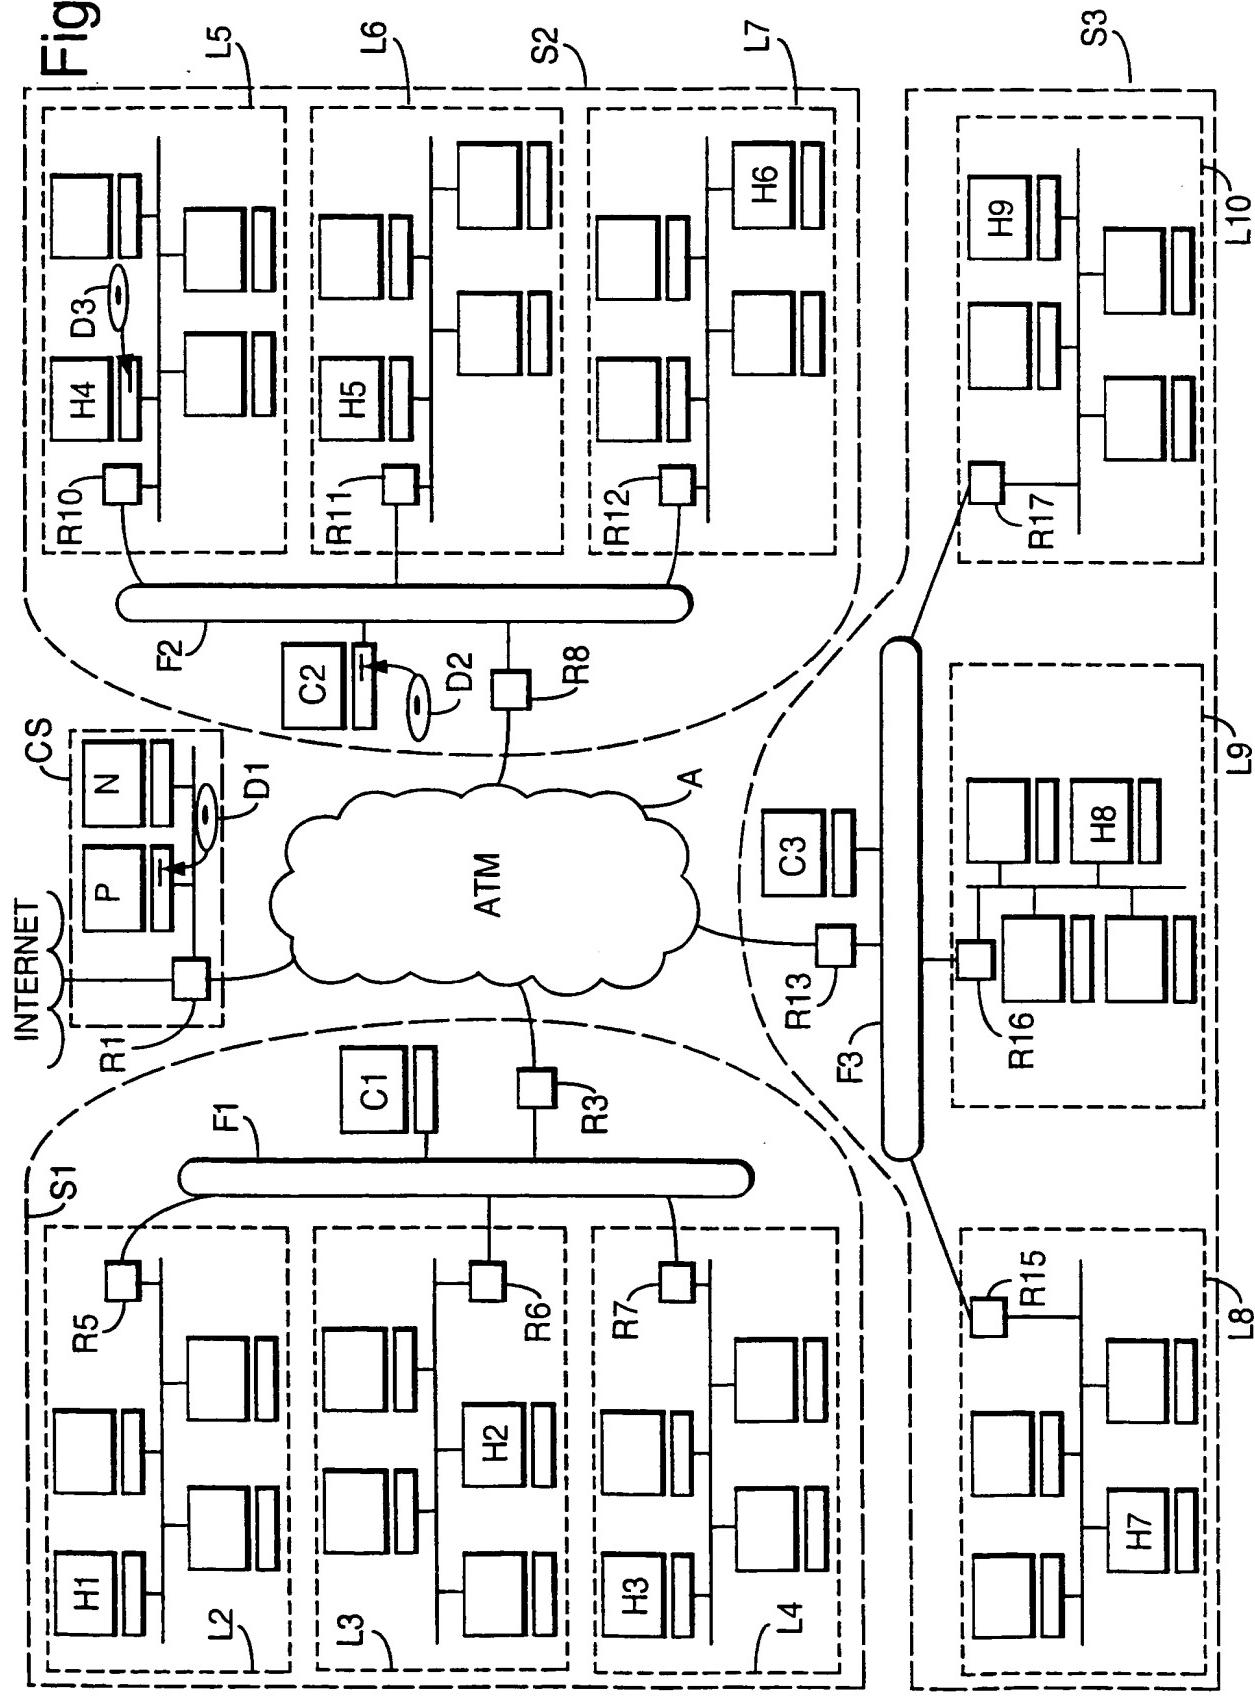


Fig.2A.

Members of Set (site level)	IEEE 802 address
H1,H2,H3	C000:0000:0001
H1,H2	C000:0000:0002
H1,H3	C000:0000:0003
H2,H3	C000:0000:0004
H1	0400:0000:0001
H2	0400:0000:0002
H3	0400:0000:0003

Fig.2B.

Members of Set	IP address
C1,C2,C3	229.274.1.27
C1,C2	229.274.1.28
C1,C3	229.274.1.29
C2,C3	229.274.1.30
C1	147.144.101.2
C2	147.144.102.2
C3	147.144.103.2

3/7

Fig.3.

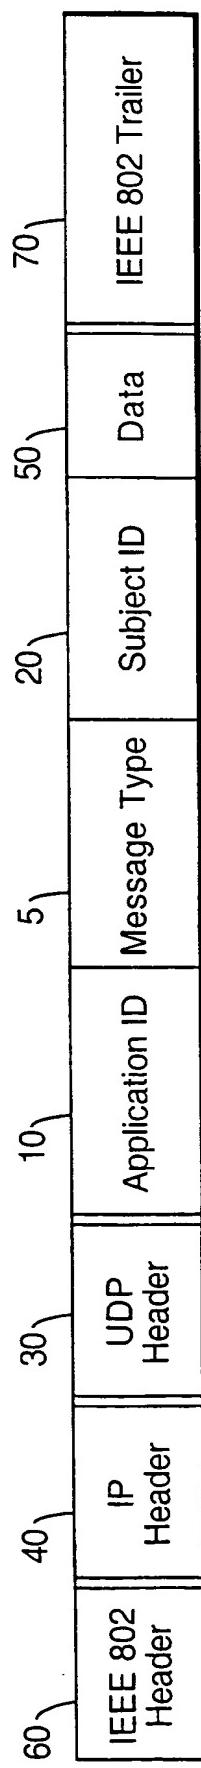


Fig.4A.

Members of Set (site level)	IEEE 802 address	Subjects subscribed to
H1,H2,H3	C000:0000:0001	none
H1,H2	C000:0000:0002	none
H1,H3	C000:0000:0003	none
H2,H3	C000:0000:0004	none
H1	0400:0000:0001	pool
H2	0400:0000:0002	none
H3	0400:0000:0003	none

Fig.4B.

Members of Set (organisation level)	IP address	Subjects subscribed to
C1,C2,C3	229.274.1.27	none
C1,C2	229.274.1.28	none
C1,C3	229.274.1.29	none
C2,C3	229.274.1.30	none
C1	147.144.101.2	pool
C2	147.144.102.2	none
C3	147.144.103.2	none

Fig.4C.

Members of Set (site level)	IEEE 802 address	Subjects subscribed to
H1,H2,H3	C000:0000:0001	none
H1,H2	C000:0000:0002	pool
H1,H3	C000:0000:0003	none
H2,H3	C000:0000:0004	none
H1	0400:0000:0001	none
H2	0400:0000:0002	none
H3	0400:0000:0003	none

Fig.5A.

Members of Set (site level)	IEEE 802 address	Subjects subscribed to
H1,H2,H3	C000:0000:0001	pool
H1,H2	C000:0000:0002	skiing
H1,H3	C000:0000:0003	none
H2,H3	C000:0000:0004	weight-lifting; wrestling
H1	0400:0000:0001	cycling; golf
H2	0400:0000:0002	rugby
H3	0400:0000:0003	hockey

6/7

Fig.5B.

Members of Set (site level)	IEEE 802 address	Subjects subscribed to
H4,H5,H6	C000:0000:0011	swimming
H4,H5	C000:0000:0012	fishing,pool
H4,H6	C000:0000:0013	golf
H5,H6	C000:0000:0014	none
H4	0400:0000:0011	tennis
H5	0400:0000:0012	basketball
H6	0400:0000:0013	hockey

Fig.5C.

Members of Set (site level)	IEEE 802 address	Subjects subscribed to
H7,H8,H9	C000:0000:0021	athletics
H7,H8	C000:0000:0022	boxing
H7,H9	C000:0000:0023	none
H8,H9	C000:0000:0024	badminton
H7	0400:0000:0021	horse-riding; cricket; squash
H8	0400:0000:0022	basketball; windsurfing; wrestling; yachting
H9	0400:0000:0023	golf; gymnastics

Fig.6.

Members of Set (organisation level)	IP address	Subjects subscribed to
C1,C2,C3	229.274.1.27	golf
C1,C2	229.274.1.28	pool,hockey
C1,C3	229.274.1.29	wrestling
C2,C3	229.274.1.30	basketball
C1	147.144.101.2	skiing, rugby, cycling, weightlifting
C2	147.144.102.2	swimming, fishing, tennis
C3	147.144.103.2	athletics, boxing, badminton, horse-riding, cricket, squash, windsurfing, yachting, gymnastics

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 99/03114

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/18

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 933 938 A (SHEEHY DAVID B) 12 June 1990 (1990-06-12) column 1, line 48 -column 2, line 4 abstract; figure 1	15
A		16, 17
X	EP 0 361 649 A (DIGITAL EQUIPMENT CORPORATION) 4 April 1990 (1990-04-04) column 7, line 49 -column 8, line 19 column 14, line 39 -column 18, line 53	15
	-/-	

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the International filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the International filing date but later than the priority date claimed

"T" later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the Invention

"X" document of particular relevance; the claimed Invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed Invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the International search

10 January 2000

Date of mailing of the International search report

28/01/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Ströbeck, A

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/03114

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>KOWALCHUK R J ET AL: "vnews: A Multicast, Multimedia News Service With Virtual Messages"</p> <p>PROCEEDINGS OF THE IEEE ANNUAL INTERNATIONAL PHOENIX CONFERENCE ON COMPUTERS AND COMMUNICATIONS, US, NEW YORK, IEEE, vol. CONF. 15, page 44-50 XP000594774</p> <p>ISBN: 0-7803-3256-3</p> <p>page 46, left-hand column, line 1 -page 48, right-hand column, line 32</p>	1,8, 12-14
A	<p>DEERING S E ET AL: "MULTICAST ROUTING IN DATAGRAM INTERNETWORKS AND EXTENDED LANS"</p> <p>ACM TRANSACTIONS ON COMPUTER SYSTEMS, vol. 8, no. 2, 1 May 1990 (1990-05-01), pages 85-110, XP000137193</p> <p>page 106, line 25 - line 37</p>	15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 99/03114

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 4933938	A	12-06-1990	NONE		
EP 0361649	A	04-04-1990	US 4864559 A AT 117146 T AU 602636 B AU 3729189 A CA 1318959 A DE 68920527 D DE 68920527 T JP 2149039 A US 5079767 A	05-09-1989 15-01-1995 18-10-1990 03-05-1990 08-06-1993 23-02-1995 07-09-1995 07-06-1990 07-01-1992	

REC'D 29 DEC 2000
WIPO PCT

5

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

(PCT Article 36 and Rule 70)

Applicant's or agent's file reference A25570 WO	FOR FURTHER ACTION See Notification of Transmittal of International Preliminary Examination Report (Form PCT/IPEA/416)	
International application No. PCT/GB99/03114	International filing date (day/month/year) 17/09/1999	Priority date (day/month/year) 18/09/1998
International Patent Classification (IPC) or national classification and IPC H04L12/18		
Applicant BRITISH TELECOMMUNICATIONS PUBLIC LIM.. ET AL.		

1. This international preliminary examination report has been prepared by this International Preliminary Examining Authority and is transmitted to the applicant according to Article 36.

2. This REPORT consists of a total of 6 sheets, including this cover sheet.

This report is also accompanied by ANNEXES, i.e. sheets of the description, claims and/or drawings which have been amended and are the basis for this report and/or sheets containing rectifications made before this Authority (see Rule 70.16 and Section 607 of the Administrative Instructions under the PCT).

These annexes consist of a total of 12 sheets.

3. This report contains indications relating to the following items:

- I Basis of the report
- II Priority
- III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- IV Lack of unity of invention
- V Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- VI Certain documents cited
- VII Certain defects in the international application
- VIII Certain observations on the international application

Date of submission of the demand 17/02/2000	Date of completion of this report 27.12.2000
Name and mailing address of the international preliminary examining authority:  European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Tx: 523656 epmu d Fax: +49 89 2399 - 4465	Authorized officer Agreda Labrador, A Telephone No. +49 89 2399 8263



**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB99/03114

I. Basis of the report

1. This report has been drawn on the basis of (substitute sheets which have been furnished to the receiving Office in response to an invitation under Article 14 are referred to in this report as "originally filed" and are not annexed to the report since they do not contain amendments (Rules 70.16 and 70.17).):

Description, pages:

8-15 as originally filed

1-7,7a as received on 27/07/2000 with letter of

26/07/2000

Claims, No.:

1-12 with telefax of 24/11/2000

Drawings, sheets:

1/7-7/7 as originally filed

2. With regard to the **language**, all the elements marked above were available or furnished to this Authority in the language in which the international application was filed, unless otherwise indicated under this item.

These elements were available or furnished to this Authority in the following language: , which is:

- the language of a translation furnished for the purposes of the international search (under Rule 23.1(b)).
 - the language of publication of the international application (under Rule 48.3(b)).
 - the language of a translation furnished for the purposes of international preliminary examination (under Rule 55.2 and/or 55.3).

3. With regard to any nucleotide and/or amino acid sequence disclosed in the international application, the international preliminary examination was carried out on the basis of the sequence listing:

- contained in the international application in written form.
 - filed together with the international application in computer readable form.
 - furnished subsequently to this Authority in written form.
 - furnished subsequently to this Authority in computer readable form.
 - The statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.
 - The statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished.

4. The amendments have resulted in the cancellation of:

INTERNATIONAL PRELIMINARY EXAMINATION REPORT

International application No. PCT/GB99/03114

the description, pages:
 the claims, Nos.: 13-17
 the drawings, sheets:

5. This report has been established as if (some of) the amendments had not been made, since they have been considered to go beyond the disclosure as filed (Rule 70.2(c)):

(Any replacement sheet containing such amendments must be referred to under item 1 and annexed to this report.)

6. Additional observations, if necessary:

V. Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	1-12
	No:	Claims	
Inventive step (IS)	Yes:	Claims	1-12
	No:	Claims	
Industrial applicability (IA)	Yes:	Claims	1-12
	No:	Claims	

2. Citations and explanations see separate sheet

VII. Certain defects in the international application

The following defects in the form or contents of the international application have been noted:
see separate sheet

VIII. Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:
see separate sheet

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT - SEPARATE SHEET**

International application No. PCT/GB99/03114

Reference is made to the following documents:

- D1: EP-A-0 854 604 (BODNER R) 22 July 1998
- D2: US-A-4 933 938 (SHEEHY DAVID B) 12 June 1990
- D3: EP-A-0 361 649 (DIGITAL EQUIPMENT CORPORATION) 4 April 1990
- D4: KOWALCHUK R J ET AL: 'vnews: A Multicast, Multimedia News Service With Virtual Messages' PROCEEDINGS OF THE IEEE ANNUAL INTERNATIONAL PHOENIX CONFERENCE ON COMPUTERS AND COMMUNICATIONS, US, NEW YORK, IEEE, vol. CONF. 15, page 44-50 XP000594774 ISBN: 0-7803-3256-3
- D5: DEERING S E ET AL: 'MULTICAST ROUTING IN DATAGRAM INTERNETWORKS AND EXTENDED LANS' ACM TRANSACTIONS ON COMPUTER SYSTEMS, vol. 8, no. 2, 1 May 1990, pages 85-110, XP000137193

Re Item V: Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. The invention concerns a method (claim 1), a corresponding apparatus (13), a program storage device (claim 11) and a computer program (claim 12) for operating a transmitter in a multicast-capable network.

A transmitter has access to one or more directories storing a plurality of lists of receiver identifiers and multicast addresses corresponding to each of said lists for use in said network. The method of the invention consists of obtaining a list of receiver identifiers corresponding to the set of recipients to which said block is to be sent and examining said one or more directories to find a multicast address corresponding to the desired list of receiver identifiers.

The invention lies in the use of a list of receiver identifiers in place of the conventional group identifier, which leads to the process of multicasting a data block being more easily adaptable to the requirements of different users.

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT - SEPARATE SHEET**

International application No. PCT/GB99/03114

For example, the directories can contain also identifiers for types of data such as news subjects accepted by each list of receivers. According to the users' requirements, each news subject can change from one list to another. In the prior art (i.e. USENET), one multicast address is assigned to each news subject resulting in a high number of multicast addresses whereas in the present invention, the number of multicast addresses is restricted to the number of lists. In other words, conventional multicast can involve a many-to-one correspondence as opposed to the one-to-one correspondence of the underlying inventive concept.

As another example of the flexibility of the invention, the lists of receivers are dynamically configurable. Lists can be created containing receivers which have not received an earlier data block.

2. D1, which was cited by the Examiner, refers to a method for use in a network which involves defining for each set of end users attached to a destination node, an **intra node multicasting address** for routing datagrams to the users addressed by the intra node multicasting address. The intra node multicasting addresses related to each set of end users are distributed to all access nodes in the network and are stored in a directory database within each access node.

Although this document discloses the use of a directory storing multicasting addresses related to each set of users, the directory does not store a plurality of lists of users (lists of receiver identifiers). Additionally, the high-speed packet switching network of D1 is not multicast-capable.

3. The rest of the documents cited in the International Search Report relate to different methods of multicasting packet networks. D2 and D3, which were relevant as to the subject-matter of claims which have been deleted from the application, refer to the subdivision of a network in two subnetworks. D4 describes a news service which provides some advantages over the current USENET system but with a similar address scheme (i.e. addresses assigned to types of data). D5 hints to the subdivision of the network in different domains so that one subdomain can be treated as a single link in a higher-level domain thus minimizing the routing costs.

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT - SEPARATE SHEET**

International application No. PCT/GB99/03114

In summary, no indication exists in any of the available prior art documents as to why a directory should contain lists of receivers and corresponding multicast addresses instead of group identifiers and corresponding group addresses. A skilled person would not adopt this solution without some indication of the benefits that it would provide. Therefore, they do not appear to be sufficiently relevant to compromise the novelty and inventiveness (Articles 33(1)-(3) PCT) of the independent claims.

4. Consequently, independent claims 1, 9, 11 and 12 appear to meet the requirements of Articles 33(1)-(3) PCT.
5. Claims 2-8 and 10 refer to advantageous embodiments of the invention and are dependent on one of the aforementioned independent claims. Therefore, they also appear to meet the requirements of Articles 33(1)-(3) PCT.

Re Item VII: Certain defects in the international application

1. The features of the claims are not provided with reference signs placed in parentheses (Rule 6.2(b) PCT).
3. The description should have been adapted to the amended set of claims (Rule 5.1(a)(iii) PCT).

Re Item VIII: Certain observations on the international application

1. Claim 1 should specify that the recipients of data blocks are computers, to be consistent (Article 6 PCT) with claim 9.
2. The category of independent claim 9 is not clear (Article 6 PCT). An apparatus cannot be defined by a set of instructions being executable to carry out a method. The apparatus should have been defined in terms of structural features.

M 21.07.00
1

MULTICASTING

The present invention relates to a method of operating a transmitter to multicast data blocks over a network.

5

In network terminology, sending a data block (i.e. a message, packet, frame or the like) across a network to a single recipient is known as unicasting. Sending a data block to all receivers connected to a network is known as broadcasting. Sending a data block to a set of recipients that form a subset of all receivers connected to a 10 network is referred to as multicasting. Such a set of recipients is known as a multicast group.

In broadcast networks (e.g. nearly all Local Area Networks), all receivers receive any data block placed the network. Each receiver has a network interface card 15 that identifies frames that are intended for that receiver. In order to identify which multicast data blocks are intended for that receiver, the network interface card must store an indication of each multicast address for which the receiver is a recipient.

20 In networks comprising switching nodes interconnected by point-to-point links, each switching node must store, for each multicast group, an indication of which of the links connected to it lead toward members of the multicast group.

In both types of networks therefore, the amount of memory required in network 25 components increases with the number of multicast addresses in use.

Conventional use of multicast networks often results in the use of more multicast addresses than is necessary. Furthermore, the amount of traffic generated in passing one or more multicast messages across the network is often unnecessarily 30 high.

European Patent Application EP 0 854 604 A1 discloses a fast packet switching network that provides interconnection for peer networks that support group addressing.

M 27.07.00
2

Although the use of multicasting is currently at a level where the above problems do not impact heavily on network performance, a rapid rise in the use of multicast is likely to take place over the next few years. This means that there is a need to
5 provide a more efficient way of using networks that provide a multicast capability.

According to a first aspect of the present invention, there is provided a method of operating a transmitter of data blocks connected via a network to a plurality of receivers of said data blocks, said method comprising:

10

finding the set of recipients to which a data block is to be sent; and

transmitting said data block over said network to said set of recipients;

15 said method being characterised by:

said network being multicast-capable;

20 said transmitter having access to one or more directories listing sets of recipients selected from said receivers and corresponding multicast addresses for use in said multicast-capable network;

examining said one or more directories to find respective one or more multicast addresses corresponding to the set of recipients found in said examining step;

25

addressing one or more data blocks to the respective multicast addresses found in said examining step.

By having one or more directories that list possible selections of recipients and
30 corresponding multicast addresses, multicast data blocks can be addressed based on the set of recipients to which it is found the data block is to be sent. This leads to the process of multicasting a data block being more easily adaptable to the requirements of different users.

M 27.07.00

For groups with only one member, the directories can provide a unicast address instead of a multicast address.

To give an example, in some embodiments of the present invention, said finding
5 step comprises:

- a) receiving one or more indications that an earlier data block addressed to a selected set of recipients was not successfully received by one or more of said recipients; and
- b) analysing said indications to find the subset of said selected set of recipients
10 that did not successfully receive said earlier data block;
wherein said examining step involves examining said one or more directories to find the multicast address corresponding to said subset of recipients.

By finding the set of recipients that require re-transmission of the data block, and
15 then finding from the one or more directories the multicast address that corresponds to that set of recipients, the re-transmission can easily be directed only towards those recipients that failed to receive the earlier data block. This results in less load being placed on the network than would occur in accordance with conventional methods where the data block is re-transmitted to the same
20 multicast address as before once one or more recipients have indicated that they have not received the data block.

To give another example of the improved adaptability provided by operating the transmitter in accordance with the present invention, in some embodiments said
25 finding step involves:

- a) determining that a general data block is to be sent to a selected plurality of sets of recipients; and
- b) unifying said plurality of sets of recipients to find a unified set of recipients;
wherein said examining step involves examining said one or more directories to
30 find from said address data the multicast address which corresponds to said unified set.

M 27.07.00
4

In this case only one multicast message need be sent across the network, as opposed to the plurality of multicast messages that would have to be sent using a conventional method.

- 5 It is to be expected that 'push' news services, Publish and Subscribe services and the like will soon migrate to the use of multicasting. 'Push' news services normally allow a user to select subjects of interest. One way of achieving the bandwidth savings offered by the use of multicast (whilst allowing a user to be selective as to which subjects he received) would be to allocate a multicast 10 address to each subject. However, given the large number of possible subjects this method would result in the components of the network having to handle a large number of multicast addresses.

- To address this problem, in a further embodiment of the present invention, said 15 transmitter further has access to type data listing data block type identifiers and corresponding sets of recipients, and said finding step involves:
- a) finding a type identifier associated with said data block; and
 - b) examining said type data to find the set of recipients associated with said type identifier.

20

The transmitter might update the type data in response to a request from a recipient to be sent data blocks containing a given type identifier, or in response to a request not to receive any further data blocks containing that type identifier.

- 25 The data block type identifier might, for example, indicate to which subject the information carried in the data block relates. In conventional methods of operating a transmitter to use a multicast-capable network, a one-to-one subject-to-multicast address mapping is used. However, it will be seen that the above embodiment enables a many-to-one subject-to-multicast address mapping to be used instead.
- 30 In networks where the number of selections of recipients is less than the number of subjects, the number of multicast addresses that must be handled by the network is reduced.

1 27.07.00

It is even possible to utilise the present invention to reduce the number of multicast addresses that must be handled by a network where the number of possible sets of recipients is very high (it will be realised that the number of sets of recipients grows exponentially with the number of receivers).

5

One way of achieving this is to divide the recipients into groups and provide a separate directory for each group. The number of multicast addresses required then grows exponentially with the number of receivers in each group, rather than growing exponentially with the number of receivers connected to the network.

10

Another way of achieving this is to provide, in accordance with a second aspect of the present invention, a method of operating an internetwork, said internetwork comprising:

a first level source of data blocks connected via a first level network to a plurality
15 of first level network receivers which include a plurality of second level
transmitters in turn connected via respective second networks to respective
pluralities of second level network receivers, said method comprising the steps of:

operating said first transmitter to transmit data blocks over said first network; and
20

operating each of said second level transmitters to forward data blocks over the
corresponding second level network

said method being characterised in that:

25

said second level transmitters have access to respective second level network
directories, each second level network directory listing data block type identifiers
and corresponding multicast addresses for the second level network to which the
directory relates; and

30 further operating said first transmitter to include a data block type identifier in each
data block;

said second level transmitters forwarding said data blocks by:

M 27.07.00
6

- a) extracting said type identifier from a data block received from the first transmitter;
- b) examining the appropriate second level network directory to find the second level network multicast address corresponding to said type identifier;
- 5 c) addressing said data block to said second level network multicast address; and
- d) transmitting said data block over said second level network.

The term internetwork is used here to mean a plurality of interconnected networks (which themselves therefore form a larger network). Here the networks are 10 arranged into a hierarchy, the connections between the first and second level networks being made by respective second level transmitters. It will be realised that the second level network directories are similar to the group directories mentioned above and result in a corresponding reduction in the number of multicast addresses that must be handled in the larger network.

15

In hierarchical embodiments of the present invention one or more of the smaller networks can be operated in accordance with one or more of the above embodiments.

- 20 In embodiments of the second aspect of the present invention, the transmitter at one level of the hierarchy (upper transmitter) places a data block type identifier in each data block that it transmits – the transmitter at the next level down in the hierarchy (the lower transmitter) is provided with access to a directory which sets out the multicast addresses associated with each data block type identifier. On 25 receiving a data block from the upper transmitter, the lower transmitter reads the data block type identifier, finds the multicast address associated with that data block type identifier, re-addresses the data block accordingly, and forwards the data block across its network. The hierarchy can have any number of levels.
- 30 In a hierarchical arrangement a request from a recipient to be sent data blocks of a particular type need only travel up the hierarchy as far as the intermediate transmitter that already receives data blocks of that type. This advantageously reduces network traffic.

M 27.07.00
7

Those skilled in the art will realise that where the smaller networks operate in accordance with the Internet Protocol suite, administrative scoping (a technique which limits the nodes through which a multicast message may pass) can be used to enable the re-use of multicast addresses in separate smaller networks, thereby 5 decreasing further the number of multicast addresses that need be handled by the larger network.

By having the transmitters intermediate the smaller networks cache data blocks for a predetermined length of time, re-transmission of the data block can be limited to 10 the smaller network in which the transmission failure occurred. Similarly, flow-control can be restricted to a smaller network that is experiencing congestion.

It will be realised that transmitters at all levels of the hierarchy can be provided with a directory listing multicast addresses for the smaller networks they transmit 15 data blocks over.

Thus, according to a third aspect of the present invention, there is provided a method of operating a packet network to transmit a plurality of packets to respective different subsets of possible recipients, said network comprising a 20 plurality of interconnected nodes, said method comprising and being characterised by the steps of, for one or more subnetworks of the network, assigning a common local multicast address for use in said subnetwork to packets which are destined for different subsets of possible recipients but which share a subset of forwarding nodes within said subnetwork.

25

According to a fourth aspect of the present invention, there is provided a data block forwarding apparatus operable to forward data blocks received from a first network over a second network to a set of second network recipients selected from a plurality of second network receivers, said apparatus comprising:

- 30 an input connectable to said first network to receive data blocks;
 an output connectable to said second network;
 one or more processors;
 a program store storing instructions executable by said one or more processors to forward the data block via said output over said second network;

M 27.07.00
7a

said data block forwarding apparatus being characterised by:

- said received data blocks containing a data block type identifier;
- 5 a forwarding database store storing data representing data block type identifiers and corresponding sets of recipients;
- a directory store storing data representing sets of recipients and corresponding multicast addresses; and
- said set of instructions being executable to forward the data block by:

10

- a) reading said data block type identifier from a data block received from said first network;
- b) examining said forwarding database store to find the set of recipients corresponding to said data block type identifier;
- 15 c) examining said directory store to find the multicast address associated with the set of recipients found in step b); and
- d) re-addressing said data block to the multicast address found in step c).

Such apparatus provides an transmitter intermediate two levels of the hierarchy
20 mentioned above.

There now follows, by way of example only, a description of specific embodiments of the present invention. The description is to be read in conjunction with the accompanying drawings, in which:

CLAIMS

1. A method of operating a transmitter to transmit a data block to a plurality of recipients selected from a plurality of receivers connected to said transmitter via a multicast-capable network, said method comprising:

finding a multicast address to which said data block is to be sent, said multicast address being suitable for use in said multicast-capable network;

10 addressing said data block to said multicast address; and

transmitting said data block over said network;

said method being characterised in that:

15

said transmitter has access to one or more directories storing:

a) a plurality of lists of receiver identifiers; and

20 b) for each of said lists, a multicast-address suitable for use in said multicast-capable network; and

said multicast address finding step comprises:

25 a) obtaining a list of receiver identifiers, said list corresponding to the set of recipients to which said data block is to be sent; and

b) examining said one or more directories to find a multicast address corresponding to the list of receiver identifiers obtained in step a).

30

2. A method according to claim 1 wherein:

said obtaining step comprises:

- a) receiving one or more indications that an earlier data block addressed to a selected set of receivers was not successfully received by one or more of said set of receivers; and
- b) analysing said indications to generate a list of receiver identifiers, each receiver identifier in said list identifying a recipient that did not successfully receive said earlier data block.

3. A method according to claim 1 wherein:
said obtaining step involves:
- 10 a) determining that a general data block is to be sent to recipients included in one or more of a selected plurality of said lists; and
b) unifying said selected plurality of lists to find a unified list of receiver identifiers.
4. A method according to claim 1 wherein said transmitter further has access to
15 type data listing data block type identifiers, and a list of recipient identifiers for each data block type identifier, wherein said obtaining step involves:
a) finding a type identifier associated with said data block; and
b) examining said type data to find a list of receiver identifiers associated with said type identifier.
- 20 5. A method according to claim 4 wherein said type identifier is a subject-matter identifier indicating the subject-matter to which the data in the data block relates.
6. A method according to claim 4 or 5 wherein said type identifier finding step
25 involves extracting a type identifier from a data block received at said transmitter.
7. A method according to claim 1 wherein said transmitter has access to a plurality of group directories for respective groups of receivers.
- 30 8. A method according to any preceding claim wherein the format of said multicast address is in accordance with the Internet Protocol suite.

9. A transmitter operable to transmit data blocks to a set of recipient computers selected from a plurality of receiver computers connectable to said transmitter computer via a multicast-capable network, said apparatus comprising:

5 an output connectable to said network;

one or more processors;

a program store storing instructions executable by said one or more processors to
10 transmit the data block via said output over said network;

said set of instructions being executable to transmit the data block by:

finding a multicast address to which said data block is to be sent, said multicast
15 address being suitable for use in said multicast-capable network;

addressing said data block to said multicast address; and

transmitting said data block over said network;

20

said transmitter being characterised by:

having access to a directory store storing:

25 a) list data representing lists of receiver identifiers; and

b) for each of said lists, a multicast address suitable for use in said multicast-capable network; and

30 said set of instructions being executable to find said multicast address by:

a) obtaining a list of receiver identifiers, said list corresponding to the set of recipients to which said data block is to be sent; and

b) examining said one or more directories to find a multicast address corresponding to the list of receiver identifiers obtained in step a).

10. A transmitter according to claim 9 wherein:

5

said transmitter further has access to type data listing data block type identifiers, and a list of recipient identifiers for each data block type identifier; and said set of instructions being further executable to obtain said list of receiver identifiers by:

- 10 a) finding a type identifier associated with said data block; and
b) examining said type data to find a list of receiver identifiers associated with said type identifier.

- 15 11. A program storage device readable by a processing apparatus, said device embodying a program of instructions executable by the processing apparatus to perform method steps for transmitting a data block over a network to a set of recipients selected from a plurality of receivers, said method steps comprising steps according to any one of claims 1 to 8.

20

12. A computer program comprising computer program code adapted to perform the method steps of any one of claims 1 to 8 when said program is executed by a computer.

25

PATENT COOPERATION TREATY

PCT

nkt

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference A25570 WO	FOR FURTHER ACTION see Notification of Transmittal of International Search Report (Form PCT/ISA/220) as well as, where applicable, item 5 below.	
International application No. PCT/GB 99/ 03114	International filing date (<i>day/month/year</i>) 17/09/1999	(Earliest) Priority Date (<i>day/month/year</i>) 18/09/1998
Applicant BRITISH TELECOMMUNICATIONS PUBLIC LIM.. ET AL.		

This International Search Report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This International Search Report consists of a total of 3 sheets.

It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

- a. With regard to the **language**, the international search was carried out on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.

- the international search was carried out on the basis of a translation of the international application furnished to this Authority (Rule 23.1(b)).
- b. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, the international search was carried out on the basis of the sequence listing :
- contained in the international application in written form.
- filed together with the international application in computer readable form.
- furnished subsequently to this Authority in written form.
- furnished subsequently to this Authority in computer readable form.
- the statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.
- the statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished

2. Certain claims were found unsearchable (See Box I).

3. Unity of invention is lacking (see Box II).

4. With regard to the **title**,

- the text is approved as submitted by the applicant.
- the text has been established by this Authority to read as follows:

HIERARCHICAL MULTICASTING

5. With regard to the **abstract**,

- the text is approved as submitted by the applicant.
- the text has been established, according to Rule 38.2(b), by this Authority as it appears in Box III. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.

6. The figure of the **drawings** to be published with the abstract is Figure No.

- as suggested by the applicant.
- because the applicant failed to suggest a figure.
- because this figure better characterizes the invention.

1 _____

None of the figures.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/03114

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L12/18

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 933 938 A (SHEEHY DAVID B) 12 June 1990 (1990-06-12) column 1, line 48 -column 2, line 4 abstract; figure 1	15
A	---	16, 17
X	EP 0 361 649 A (DIGITAL EQUIPMENT CORPORATION) 4 April 1990 (1990-04-04) column 7, line 49 -column 8, line 19 column 14, line 39 -column 18, line 53 ----	15
	-/-	

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

° Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

10 January 2000

Date of mailing of the international search report

28/01/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ströbeck, A

INTERNATIONAL SEARCH REPORT

International Application No

GB 99/03114

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	KOWALCHUK R J ET AL: "vnews: A Multicast, Multimedia News Service With Virtual Messages" PROCEEDINGS OF THE IEEE ANNUAL INTERNATIONAL PHOENIX CONFERENCE ON COMPUTERS AND COMMUNICATIONS, US, NEW YORK, IEEE, vol. CONF. 15, page 44-50 XP000594774 ISBN: 0-7803-3256-3 page 46, left-hand column, line 1 -page 48, right-hand column, line 32 ---	1,8, 12-14
A	DEERING S E ET AL: "MULTICAST ROUTING IN DATAGRAM INTERNETWORKS AND EXTENDED LANS" ACM TRANSACTIONS ON COMPUTER SYSTEMS, vol. 8, no. 2, 1 May 1990 (1990-05-01), pages 85-110, XP000137193 page 106, line 25 - line 37 -----	15

INTERNATIONAL SEARCH REPORT

International Application No

T/GB 99/03114

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 4933938	A 12-06-1990	NONE		
EP 0361649	A 04-04-1990	US 4864559 A	05-09-1989	
		AT 117146 T	15-01-1995	
		AU 602636 B	18-10-1990	
		AU 3729189 A	03-05-1990	
		CA 1318959 A	08-06-1993	
		DE 68920527 D	23-02-1995	
		DE 68920527 T	07-09-1995	
		JP 2149039 A	07-06-1990	
		US 5079767 A	07-01-1992	

8/PRTS

1

MULTICASTING

The present invention relates to a method of operating a transmitter to multicast data blocks over a network.

5

In network terminology, sending a data block (i.e. a message, packet, frame or the like) across a network to a single recipient is known as unicasting. Sending a data block to all receivers connected to a network is known as broadcasting. Sending a data block to a set of recipients that form a subset of all receivers connected to a 10 network is referred to as multicasting. Such a set of recipients is known as a multicast group.

In broadcast networks (e.g. nearly all Local Area Networks), all receivers receive any data block placed the network. Each receiver has a network interface card 15 that identifies frames that are intended for that receiver. In order to identify which multicast data blocks are intended for that receiver, the network interface card must store an indication of each multicast address for which the receiver is a recipient.

20 In networks comprising switching nodes interconnected by point-to-point links, each switching node must store, for each multicast group, an indication of which of the links connected to it lead toward members of the multicast group.

25 In both types of networks therefore, the amount of memory required in network components increases with the number of multicast addresses in use.

Conventional use of multicast networks often results in the use of more multicast addresses than is necessary. Furthermore, the amount of traffic generated in passing one or more multicast messages across the network is often unnecessarily 30 high.

Although the use of multicasting is currently at a level where these problems do not impact heavily on network performance, a rapid rise in the use of multicast is

likely to take place over the next few years. This means that there is a need to provide a more efficient way of using networks that provide a multicast capability.

According to a first aspect of the present invention, there is provided a method of
5 operating a transmitter of data blocks connected via a network having a multicast capability to a plurality of receivers of said data blocks, said transmitter having access to one or more directories listing sets of recipients selected from said receivers and corresponding multicast addresses; said method comprising:

- 10 a) finding the set of recipients to which a data block is to be sent;
 - b) examining said one or more directories to find one or more respective multicast addresses corresponding to said set of recipients;
 - 15 c) addressing one or more data blocks to said one or more respective multicast addresses; and
 - d) transmitting said one or more data blocks over said network.
- 20 By having one or more directories that list possible selections of recipients and corresponding multicast addresses, multicast data blocks can be addressed based on the set of recipients to which it is found the data block is to be sent. This leads to the process of multicasting a data block being more easily adaptable to the requirements of different users.

25

For groups with only one member, the directories can provide a unicast address instead of a multicast address.

- To give an example, in some embodiments of the present invention, said finding
30 step comprises:
 - a) receiving one or more indications that an earlier data block addressed to a selected set of recipients was not successfully received by one or more of said recipient's; and

b) analysing said indications to find the subset of said selected set of recipients that did not successfully receive said earlier data block;
wherein said examining step involves examining said one or more directories to find the multicast address corresponding to said subset of recipients.

5

By finding the set of recipients that require re-transmission of the data block, and then finding from the one or more directories the multicast address that corresponds to that set of recipients, the re-transmission can easily be directed only towards those recipients that failed to receive the earlier data block. This 10 results in less load being placed on the network than would occur in accordance with conventional methods where the data block is re-transmitted to the same multicast address as before once one or more recipients have indicated that they have not received the data block.

15 To give another example of the improved adaptability provided by operating the transmitter in accordance with the present invention, in some embodiments said finding step involves:

a) determining that a general data block is to be sent to a selected plurality of sets of recipients; and
20 b) unifying said plurality of sets of recipients to find a unified set of recipients;
wherein said examining step involves examining said one or more directories to find from said address data the multicast address which corresponds to said unified set.

25 In this case only one multicast message need be sent across the network, as opposed to the plurality of multicast messages that would have to be sent using a conventional method.

It is to be expected that 'push' news services, Publish and Subscribe services and 30 the like will soon migrate to the use of multicasting. 'Push' news services normally allow a user to select subjects of interest. One way of achieving the bandwidth savings offered by the use of multicast (whilst allowing a user to be selective as to which subjects he received) would be to allocate a multicast address to each subject. However, given the large number of possible subjects

this method would result in the components of the network having to handle a large number of multicast addresses.

To address this problem, in a further embodiment of the present invention, said
5 transmitter further has access to type data listing data block type identifiers and corresponding sets of recipients, and said finding step involves:
a) finding a type identifier associated with said data block; and
b) examining said type data to find the set of recipients associated with said type identifier.

10

The transmitter might update the type data in response to a request from a recipient to be sent data blocks containing a given type identifier, or in response to a request not to receive any further data blocks containing that type identifier.

15 The data block type identifier might, for example, indicate to which subject the information carried in the data block relates. In conventional methods of operating a transmitter to use a multicast-capable network, a one-to-one subject-to-multicast address mapping is used. However, it will be seen that the above embodiment enables a many-to-one subject-to-multicast address mapping to be used instead.
20 In networks where the number of selections of recipients is less than the number of subjects, the number of multicast addresses that must be handled by the network is reduced.

It is even possible to utilise the present invention to reduce the number of
25 multicast addresses that must be handled by a network where the number of possible sets of recipients is very high (it will be realised that the number of sets of recipients grows exponentially with the number of receivers).

One way of achieving this is to divide the recipients into groups and provide a
30 separate directory for each group. The number of multicast addresses required then grows exponentially with the number of receivers in each group, rather than growing exponentially with the number of receivers connected to the network.

Another way of achieving this is to provide, in accordance with a second aspect of the present invention, a method of operating an internetwork, said internetwork comprising:

- a first level transmitter of data blocks connected via a first level network to a plurality of first level network receivers which include a plurality of second level transmitters in turn connected via respective second networks to respective pluralities of second level network receivers,
- 5 said second level transmitters having access to respective second level network directories, each second level network directory listing data block type identifiers and corresponding multicast addresses for the second level network to which the directory relates, said method comprising the steps of:
 - operating said first transmitter to transmit data blocks over said first network, and further operating said first transmitter to include a data block type identifier in each data block;
 - 10 operating each of said second level transmitters to forward data blocks over the corresponding second level network by:
 - a) extracting said type identifier from a data block received from the first transmitter;
 - b) examining the appropriate second level network directory to find the second level network multicast address corresponding to said type identifier;
 - 20 c) addressing said data block to said second level network multicast address; and
 - d) transmitting said data block over said second level network.

The term internetwork is used here to mean a plurality of interconnected networks (which themselves therefore form a larger network). Here the networks are arranged into a hierarchy, the connections between the first and second level networks being made by respective second level transmitters. It will be realised that the second level network directories are similar to the group directories mentioned above and result in a corresponding reduction in the number of 30 multicast addresses that must be handled in the larger network.

In hierarchical embodiments of the present invention one or more of the smaller networks can be operated in accordance with one or more of the above embodiments.

In embodiments of the second aspect of the present invention, the transmitter at one level of the hierarchy (upper transmitter) places a data block type identifier in each data block that it transmits – the transmitter at the next level down in the 5 hierarchy (the lower transmitter) is provided with access to a directory which sets out the multicast addresses associated with each data block type identifier. On receiving a data block from the upper transmitter, the lower transmitter reads the data block type identifier, finds the multicast address associated with that data block type identifier, re-addresses the data block accordingly, and forwards the 10 data block across its network. The hierarchy can have any number of levels.

In a hierarchical arrangement a request from a recipient to be sent data blocks of a particular type need only travel up the hierarchy as far as the intermediate transmitter that already receives data blocks of that type. This advantageously 15 reduces network traffic.

Those skilled in the art will realise that where the smaller networks operate in accordance with the Internet Protocol suite, administrative scoping (a technique which limits the nodes through which a multicast message may pass) can be used 20 to enable the re-use of multicast addresses in separate smaller networks, thereby decreasing further the number of multicast addresses that need be handled by the larger network.

By having the transmitters intermediate the smaller networks cache data blocks for 25 a predetermined length of time, re-transmission of the data block can be limited to the smaller network in which the transmission failure occurred. Similarly, flow-control can be restricted to a smaller network that is experiencing congestion.

It will be realised that transmitters at all levels of the hierarchy can be provided 30 with a directory listing multicast addresses for the smaller networks they transmit data blocks over.

Thus, according to a third aspect of the present invention, there is provided a method of operating a packet network to transmit a plurality of packets to

respective different subsets of possible recipients, said method comprising the steps of, for one or more portions of the network, assigning a common local multicast address to packets which are destined for different subsets of possible recipients but which share a subset of forwarding nodes within said portion.

5

According to a fourth aspect of the present invention, there is provided a data block forwarding apparatus operable to forward data blocks received from a first network over a second network to a set of second network recipients selected from a plurality of second network receivers, said apparatus comprising:

10 an input connectable to said first network to receive data blocks containing a data block type identifier;

 an output connectable to said second network;

 one or more processors;

15 a forwarding database store storing data representing data block type identifiers and corresponding sets of recipients;

 a directory store storing data representing sets of recipients and corresponding multicast addresses;

 a program store storing instructions executable by said one or more processors to:

20 a) read said data block type identifier from a data block received from said first network;

 b) examine said forwarding database store to find the set of recipients corresponding to said data block type identifier;

25 c) examine said directory store to find the multicast address associated with the set of recipients found in step b);

 d) re-address said data block to the multicast address found in step c); and

 e) forward the data block via said output over said second network.

Such apparatus provides an transmitter intermediate two levels of the hierarchy
30 mentioned above.

There now follows, by way of example only, a description of specific embodiments of the present invention. The description is to be read in conjunction with the accompanying drawings, in which:

CLAIMS

1. A method of operating a transmitter of data blocks connected via a
5 multicast-capable network to a plurality of receivers of said data blocks, said
transmitter having access to one or more directories listing sets of recipients
selected from said receivers and corresponding multicast addresses; said method
comprising:

- 10 a) finding the set of recipients to which a data block is to be sent;

b) examining said one or more directories to find respective one or more multicast
addresses corresponding to the set of recipients found in step a);

15 c) addressing one or more data blocks to the respective multicast addresses found
in step b); and

d) transmitting said one or more data blocks over said network.

20

2. A method according to claim 1 wherein:
said finding step comprises:
a) receiving one or more indications that an earlier data block addressed to a
selected set of recipients was not successfully received by one or more of said
25 recipients; and
b) analysing said indications to find the subset of said selected set of recipients
that did not successfully receive said earlier data block;
said examining step involves examining said one or more directories to find the one
or more multicast addresses corresponding to said subset of recipients;
30 whereby said transmitting step involves re-sending said earlier data block to said
subset of recipients.

3. A method according to claim 1 wherein:

said finding step involves:

- a) determining that a general data block is to be sent to a selected plurality of sets of recipients; and
 - b) unifying said plurality of sets of recipients to find a unified set of recipients;
- 5 wherein said examining step involves examining said one or more directories to find the one or more multicast addresses which correspond to said unified set.

4. A method according to claim 1 wherein said transmitter further has access
10 to type data listing data block type identifiers and corresponding sets of recipients,
wherein said finding step involves:

- a) finding a type identifier associated with said data block; and
- b) examining said type data to find the set of recipients associated with said type identifier.

15

5. A method according to claim 4 wherein said type identifier is a subject-matter identifier indicating the subject-matter to which the data in the data block relates.

20 6. A method according to claim 4 or 5 wherein said type identifier finding step involves extracting a type identifier from a data block received at said transmitter.

7. A method according to claim 1 wherein said transmitter has access to a
25 plurality of group directories for respective groups of receivers.

8. A method of operating an internetwork, said internetwork comprising:
a first level transmitter of data blocks connected via a first level network to a plurality of first level network receivers which include a plurality of second level
30 transmitters in turn connected via respective second networks to respective pluralities of second level network receivers,
said second level transmitters having access to respective second level network directoriés, each second level network directory listing data block type identifiers

and corresponding multicast addresses for the second level network to which the directory relates, said method comprising the steps of:

operating said first transmitter to transmit data blocks over said first network, and further operating said first transmitter to include a data block type identifier in each

5 data block;

operating each of said second level transmitters to forward data blocks over the corresponding second level network by:

a) extracting said type identifier from a data block received from the first transmitter;

10 b) examining the appropriate second level network directory to find the second level network multicast address corresponding to said type identifier;

c) addressing said data block to said second level network multicast address; and

d) transmitting said data block over said second level network.

15 9. A method according to claim 8 wherein said second level transmitters cache said data blocks and operate to re-transmit data blocks over said second level network in accordance with the method of claim 2.

10. A method according to claim 8 wherein said first level transmitter has
20 access to a first level network directory listing sets of recipients and corresponding multicast addresses for the first level network and operates in accordance with the method of any one of claims 1 to 6.

11. A method according to any preceding claim wherein the format of said
25 multicast address is in accordance with the Internet Protocol suite.

12. A data block forwarding apparatus operable to forward data blocks received from a first network over a second network to a set of second network recipients selected from a plurality of second network receivers, said apparatus
30 comprising:

an input connectable to said first network to receive data blocks containing a data block type identifier;

'an output connectable to said second network;

one or more processors;

a forwarding database store storing data representing data block type identifiers and corresponding sets of recipients;

a directory store storing data representing sets of recipients and corresponding multicast addresses;

5 a program store storing instructions executable by said one or more processors to:

a) read said data block type identifier from a data block received from said first network;

b) examine said forwarding database store to find the set of recipients 10 corresponding to said data block type identifier;

c) examine said directory store to find the multicast address associated with the set of recipients found in step b);

d) re-address said data block to the multicast address found in step c); and

e) forward the data block via said output over said second network.

15

13. A program storage device readable by a processing apparatus, said device embodying a program of instructions executable by the processing apparatus to perform method steps for transmitting a data block over a network to a set of recipients selected from a plurality of network receivers, said method steps 20 comprising steps according to any one of claims 1 to 6.

14. A computer program comprising computer program code adapted to perform the method steps of any one of claims 1 to 7 when said program is executed by a computer.

25

15. A method of operating a packet network to transmit a plurality of packets to respective different subsets of possible recipients, said method comprising the steps of, for one or more portions of the network, assigning a common local multicast address to packets which are destined for different subsets of possible 30 recipients but which share a subset of forwarding nodes within said portion.

16. A method according to claim 15 wherein said assignment step involves:

extracting, from the payload of each packet, destination data identifying the subset of possible recipients to which the packet is to be sent; and

deriving an associated local multicast address on the basis of said destination data; and

transmitting said packet to said associated local multicast address.

- 5 17. A method according to claim 16 wherein said destination data comprises subject identifying data.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.